

APPLICATION
FOR
UNITED STATES LETTERS PATENT

APPLICANT NAME: Quesenberry et al.

TITLE: AUTOMATIC DATA INTERPRETATION AND
IMPLEMENTATION USING PERFORMANCE CAPACITY
MANAGEMENT FRAMEWORK OVER MANY SERVERS

DOCKET NO.: FIS920010174US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

CERTIFICATE OF MAILING UNDER 37 CFR 1.10

I hereby certify that, on the date shown below, this correspondence
is being deposited with the United States Postal Service in an
envelope addressed to the U.S. Patent & Trademark Office
P.O. Box 2327** as "Express Mail Post Office to Addressee"

Mailing Label No. ET13402212705

on 12/17/01

K.C. INQ-MARS

Name of person mailing paper

K.C. INQ-MARS 12/17/01

Signature

Date

** Arlington, VA 22202

ibmf100347000pat

**AUTOMATIC DATA INTERPRETATION AND IMPLEMENTATION
USING PERFORMANCE CAPACITY MANAGEMENT FRAMEWORK
OVER MANY SERVERS**

5

Background Of The Invention

1. Field of the Invention

The present invention relates generally to computer systems and, in particular, to a method of adding critical hardware resource capacity to a networked system of computers.

10 2. Description of Related Art

As described in U.S. Patent No. 6,148,335, the disclosure of which is incorporated herein by reference, a generalized client-server computing network has a plurality of servers and which are interconnected, either directly to each other or indirectly through one of the other servers. Each server is essentially a stand-alone
15 computer system (having one or more processors, memory devices, and communications devices), but has been adapted (programmed) for the primary purpose of providing information to individual users at a plurality of workstation clients in communication with each server. A client is a member of a class or group of computers or computer systems that uses the services of another class or group to
20 which it is not related. As used herein, "client" generally refers to any multi-purpose or limited-purpose computer adapted for use by a single individual, regardless of the manufacturer, hardware platform, operating system, and the like. The information provided by a server can be in the form of programs which run locally on a given client, or in the form of data such as files used by other programs.

25 Such networks may communicate via the Internet using conventional protocols and services which allow the transfer of various types of information, including electronic mail, simple file transfers via FTP, remote computing via TELNET, gopher searching, Usenet newsgroups, and hypertext file delivery and multimedia streaming

via the World Wide Web (WWW). A given server can be dedicated to performing one of these operations, or running multiple services. The '335 patent discloses the monitoring of server performance in a network like the Internet, and generating reports detailing performance statistics (daily, weekly, or monthly) for various server resources. Statistical parameters may include for example the number of observations; CPU utilization; system usage percentage; user usage percentage; percentage of time I/O wait is greater than some pre-selected level; run queue length; active virtual memory (AVM); free space (FRE); percentage of time CPU utilization is greater than some pre-selected level; percentage of time run queue is greater than some pre-selected level; percentage of time storage usage is greater than some pre-selected level and percentage of time paging rate is greater than some number of pages per second. Links may be provided to view additional, detailed information regarding, for example, a specific resource on a particular server. Notwithstanding the advantages of the invention of the '335 patent, there is no method or system which may act on the performance information generated on the client-server computer network to improve the performance and reliability of the network.

Summary of the Invention

Bearing in mind the problems and deficiencies of the prior art, it is therefore an object of the present invention to provide a method of monitoring and controlling network computer resources.

It is another object of the present invention to provide a method of managing computing resources on a network to improve the network's performance and reliability.

A further object of the invention is to provide a method by which critical hardware resource conditions may be addressed without direct user input.

Still other objects and advantages of the invention will in part be obvious and will in part be apparent from the specification.

The above and other objects and advantages, which will be apparent to one of skill in the art, are achieved in the present invention which is directed to, in a first aspect, an automated method of managing computing resources having a workload of a given type. The method comprises providing resource data collectors for collecting data regarding performance of the resources, in accordance with the type of workload; developing a forecast of utilization of the resources, based on historical performance data; and collecting real-time performance data regarding the resources running under the workload. The method then includes analyzing the performance data and the forecast to identify a critical resource and automatically adjusting a capacity of the resource to provide steady-state performance of the resource under the workload.

In another aspect, the present invention is directed to a program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform an automated method of managing computing resources having a workload of a given type, using resource data collectors for collecting data regarding performance of the resources in accordance with the type of workload, and a forecast of utilization of the resources based on historical performance data. The method steps comprise collecting real-time performance data regarding the resources running under the workload, analyzing the performance data and the forecast to identify a critical resource, and automatically adjusting a capacity of the resource to provide steady-state performance of the resource under the workload.

In both of the above aspects of the invention, the resources preferably comprise a server network. The method may further comprise setting threshold values for the performance data and identifying the resource in accordance with the threshold values. The method may also comprise notifying a user of the computing resources when the critical resource is a hardware resource, and notifying the user when the capacity of the hardware resource is adjusted. Preferably, the method further includes initially providing additional hardware resources available to, but un-

used by, the computing resources. Such additional hardware resources may be selected from the group consisting of CPUs, computer memory and computer disk storage.

5 In yet another aspect, the present invention is directed to a computer program product for performing an automated method of managing computing resources having a workload of a given type, using resource data collectors for collecting data regarding performance of the resources in accordance with the type of workload, and a forecast of utilization of the resources based on historical performance data. The computer program product has computer-readable program code for collecting real-
10 time performance data regarding the resources running under the workload, computer-readable program code for analyzing the performance data and the forecast to identify a critical resource, and computer-readable program code for automatically adjusting a capacity of the resource to provide steady-state performance of the resource under the workload.

15 As before, in this aspect the resources preferably comprise a server network. There may be initially provided additional hardware resources available to, but unused by, the computing resources. Such additional hardware resources may be selected from the group consisting of CPUs, computer memory and computer disk storage. The computer program product may further comprise computer-readable program
20 code for setting threshold values for the performance data and computer-readable program code for identifying the resource in accordance with the threshold values. The computer program product may also comprise computer-readable program code for notifying a user of the computing resources when the critical resource is a hardware resource, and computer-readable program code for notifying the user when
25 the capacity of the hardware resource is adjusted.

Brief Description of the Drawings

The features of the invention believed to be novel and the elements characteristic of the invention are set forth with particularity in the appended claims. The figures are for illustration purposes only and are not drawn to scale. The

invention itself, however, both as to organization and method of operation, may best be understood by reference to the detailed description which follows taken in conjunction with the accompanying drawings in which:

Fig. 1 is a flow chart of a portion of the preferred method of practicing the present invention.

Fig. 2 is a continuation of the flow chart of Fig. 1 showing the preferred method of practicing the present invention.

Fig. 3 is a schematic of one embodiment of a computer client/server network employing the method of the present invention.

Fig. 4 is a schematic of a preferred RISC computer having hardware resources which may be repartitioned among different partitioned servers.

Description of the Preferred Embodiment(s)

In describing the preferred embodiment of the present invention, reference will be made herein to Figs. 1-4 of the drawings in which like numerals refer to like features of the invention. Features of the invention are not necessarily shown to scale in the drawings.

The present invention is particularly useful in connection with the successful server resource management (SRM) methodology defined in the aforementioned U.S. Patent 6,148,335, whereby server resources are measured across multiple platforms and server trends reported by enterprise and/or server-level drill-down navigation using red/yellow/green report presentation. An online "red action list" of action plan and status is also reported. The automatic data interpretation of the present invention adds a layer of benefit by implementing a set of automatic actions based on predefined correlation algorithms. When managing hundreds of installed machines, associated support costs are reduced through use of this management automation and alert methodology. In general, the present invention takes available server resource metrics for hardware resources such as central processing unit (CPU), memory and disk storage and develop framework to automatically determine a set of actions based on measured conditions. This invention forms a closed loop whereby data is not only

collected and reported at face-value, but also enables a set of recommendations or actions to be taken against the available data, saving analysis labor and intervention. The present invention provides a method to use capacity on demand to add capacity automatically to the computer system, and to notify the user, e.g., the system manager or system analyst, when hardware capacity is added.

The present invention expands on the '335 patent and expands the list of actions to automatically recommend or implement capacity planning alternatives, such that the primary focus is server capacity planning. The present invention interprets server metrics and workload resource data across platforms, and is not limited to mainframe data; it automatically determines a set of actions based on measured conditions, and uses statistical data and deduction techniques to perform the automation. The method and system of present invention are particularly directed to monitoring and analyzing server management data, as opposed to the business data on the server. The present invention uses expected resources metric feeds and supports systems management of servers and/or Information Technology (I/T) machines; automatic interpretation is performed on the expected (server historical) data and rules set implemented; and uses relational database to archive the server history.

The method of the present invention may be described in detail in connection with the flowcharts shown in Figs. 1 and 2. Initially, in step 110, server resource collectors are installed in the server system to collect data regarding performance of the server resources and threshold values for the performance data are defined. In subsequent steps, the system begins collecting and logging steady state server metrics by first starting the data collection process 120, determining the running workload of the server system 130, starting the data collection for each workload 140, and setting the collector threshold based on the workload mix 150. The workloads and workload mix may be any combination of system and/or application processes such as web hosting, database hosting, file serving, security checking, batch processing, financial systems, network management, systems management, numerical and statistical

analysis, online processing, and the like. Using the historical (as opposed to real-time) steady state server data measurements as input 160, the method then develops a metrics associated with CPU usage and thresholds to determine the need for additional CPU capacity.

5 Based on the information previously collected and the forecast computed, the method then determines whether there are any response or resource bottlenecks 190. These are determined using specific platform metrics, such as page rate, run queue, scan rate, out-and-ready, swap rate; I/O rates, disk utilization, and the like. The server system response time is measured and statistics are correlated 200 to determine the threshold values to be set for use of the hardware resources. For example, disk storage capacity threshold values may be set at some percentage of available disk space, or CPU usage may be set at some percentage of maximum usage. If threshold exceptions are found, then the method determines whether hardware resources are an impact 210. If no hardware contention correlation is found, then the customer is notified of response time threshold exceptions 220, with no hardware issues detected. 15 If the critical resource capacity is available on-demand 230, then such capacity is adjusted, i.e., added dynamically to the server 250, and the customer is notified of action taken 260. If a hardware resource contention is detected, i.e., a critical resource, and no additional hardware capacity is available, then the customer is notified of need for capacity 240. 20 If no response time or hardware resource exceptions are detected, then regular steady state conditions 270 continue and the process continues again: (a) log date; (b) analyze and correlate data; (c) activate automation policy or alerting, as necessary; (d) notify customer of actions or resource status; and (e) continue steady state.

25 The critical hardware resource capacity added in step 250 may be any central processing unit (CPU) microprocessor, computer memory, storage, or other hardware resource necessary to maintain the system at steady state operation. Memory devices may include random access memory (RAM), read only memory (ROM), and nonvolatile memory (e.g., EPROM, flash memory, or battery-pack CMOS RAM).

Storage includes disk such as optical (e.g., CD-ROM) or magnetic drives, or other storage media. Fig. 3 depicts one possible hardware arrangement for use with the present invention. Servers 10, 12, 14, 16 are linked to each other as a plurality of network nodes operating on the same or different platforms. Each server is linked to a plurality of client computers, i.e., server 10 is linked to clients 20a, 20b, 20c, server 12 is linked to clients 22a, 22b, 22c, server 14 is linked to clients 24a, 24b, 24c, and server 16 is linked to clients 26a, 26b, 26c. The clients may be stand-alone personal computers or limited-use network computers. The links between the various clients and servers are sufficient to transfer the types of information used on the particular network, such as the aforescribed Internet protocols and services.

The server resource management (SRM) architecture collects data using a remote command facility (RCF) program on server 30 which works by executing UNIX commands to gather utilization data from one or more servers, such as by the scripting language known as PERL (practical extraction and report language) to issue the commands which gather the bulk of the data. The UNIX or other machine- or computer-readable program code used by the RCF may be stored on any of the storage media described above. The RCF process can use the low- impact "sockets" interface, and be extensible for executing data gathering commands on other brands of UNIX. RCF collects key server resource data including current CPU utilization, memory availability, I/O usage, and permanent storage (disk) capacity. An output file is generated containing the collected information, which can be stored locally on a hard disk drive or at a remote location, preferably not one of the servers being monitored. The RCF can provide a user interface for data collection by using conventional communications software such as a web browser that is adapted to display a page having commands or tool bars used to manage data collection. Other communications software can be used besides standard web browsers, such as those described in the '335 patent. If a server does not run a UNIX-type platform, other commands can be used to collect the data, such as those also described in the '335 patent. In the foregoing manner, key performance and capacity data from a wide

variety of servers becomes web-accessible. Data collection from different servers can occur at different times, i.e., there is no need for data processing system on server 30 to be continuously connected to each of the servers.

Once the data has been collected, it can be deposited into an appropriate database as described, and optionally merged with other historical data previously collected. The collected data can then be forwarded (e.g., via FTP) to a node running an analysis program, such as the Statistical Analysis System (SAS). This software provides a programming language used to analyze data processing. The analyzed data can be presented in a variety of media or formats. In one implementation, a web browser can again be used to view the analysis, by creating an HTML file which is then placed on the network (e.g., the World Wide Web) in such a manner as to be accessible and usable by the end-user.

Going beyond merely reporting server performance, the present invention has the ability to adjust automatically capacity of the hardware resources identified as deficient by the SRM. The method of the present invention, as described in Figs. 1 and 2, is programmed by conventional programming code in the RCF program. As shown in Fig. 3, the additional hardware resources, such as CPU, memory or disk storage 40, 42, is available to be linked to the server on command of the RCF. These additional hardware resources may be available to a plurality of servers, such as resource 42 is available to servers 12, 14, 16, or the additional hardware resources may be available to a single server, such as resource 42 is available only to servers 10. The links from the additional hardware resources to the servers are activated and connected on command by the RCF when added capacity is determined by the RCF to be needed by one or more of the servers. Thus, the method and system of present invention are able to analyze not only the utilization of the entirety of the server install base, but also exceptions at individual, single servers which require additional resources.

A preferred computer system on which to use the method of the present invention is a mid-level computer partitioned to operate as a plurality of separate

5 servers, and capable of being re-partitioned to reallocate critical hardware resources among the different servers. This system is depicted in schematic in Fig. 4 wherein the total computer system hardware and software resources 50 are subdivided into subsystems, here showing some of such subsystems as virtual, separate servers 10a, 10b, 10c, 10d, 10e, 10f, 10g and 10h, each linked to each other for conventional network communication. The desired hardware resources which are available to be dedicated to or shared among the virtual servers, such as CPU, memory and disk storage, are shown as 60. Initially, a predetermined amount of the CPU, memory, disk storage and other hardware resource is dedicated to each virtual server. As the network operates normally under its workload, the RCF determines whether additional resource capacity is needed at one or more of the virtual servers, in accordance with the method of the present invention. The RCF then repartitions available resources 60 to add the identified critical resource to the virtual server requiring it, to restore steady state performance. A computer system having such partitionable hardware resources is under development as a RISC 6000 system from IBM Corporation, Armonk, NY.

20 The present invention automatically determines what workload is running on the computer, starts collectors based on type of workload, sets thresholds for metrics based on workload mix, determines when metrics exceed threshold (both current and projected workload), and correlates metrics to determine if hardware capacity is the cause of the problem. Additionally, the present invention automatically correlates server metrics with available middleware metrics to enable problem detection. By this method, it is possible to determine automatically if extra capacity exists, determine resource bottlenecks using historical data, add capacity if available, automatically notify people of actions taken, and provide a customer interface to set custom resource thresholds.

While the present invention has been particularly described, in conjunction with a specific preferred embodiment, it is evident that many alternatives, modifications and variations will be apparent to those skilled in the art in light of the foregoing description. It is therefore contemplated that the appended claims will
5 embrace any such alternatives, modifications and variations as falling within the true scope and spirit of the present invention.

We claim: